arXiv:1603.08213v1 [cs.LO] 27 Mar 2016

Generating reversible circuits from higher-order functional programs

Benoît Valiron

March 20, 2016

Abstract

Boolean reversible circuits are boolean circuits made of reversible elementary gates. Despite their constrained form, they can simulate any boolean function. The synthesis and validation of a reversible circuit simulating a given function is a difficult problem. In 1973, Bennett proposed to generate reversible circuits from traces of execution of Turing machines. In this paper, we propose a novel presentation of this approach, adapted to higher-order programs. Starting with a PCF-like language, we use a monadic representation of the trace of execution to turn a regular boolean program into a circuit-generating code. We show that a circuit traced out of a program computes the same boolean function as the original program. This technique has been successfully applied to generate large oracles with the quantum programming language Quipper.

1 Introduction

Reversible circuits are linear, boolean circuits with no loops, whose elementary gates are reversible. In quantum computation, reversible circuits are mostly used as oracle: the *description* of the problem to solve. Usually, this description is given as a classical, conventional algorithm: the graph to explore [17], the matrix coefficients to process [13], *etc*. These algorithms use arbitrarily complex structures, and if some are rather simple, for example [28], others are quite complicated and make use of analytic functions [13], memory registers [2] (which thus have to be simulated), *etc*.

This paper[1] is concerned with the design of reversible circuits as *operational semantics* of a higher-order purely functional programming language. The language is expressive enough to encode most algorithms: it features recursion, pairs, booleans and lists, and it can easily be extended with additional structures if needed. This operational semantics can be understood as the *compilation* of a program into a reversible circuit.

Compiling a program into a reversible circuit is fundamentally different from compiling on a regular back-end: there is no notion of "loop", no real control flow, and all branches will be explored during the execution. In essence, a reversible circuit is the *trace* of all possible executions of a given program. Constructing a reversible circuit

[1] A shorter preprint has been accepted for publication in the Proceedings of *Reversible Computation 2016*. The final publication is available at `http://link.springer.com`.

out of the trace of execution of a program is what Bennett [3] proposed in 1973, using Turing machines. In this paper, we refer to it as Landauer embeddings [16].

In this paper, we build up on this idea of circuit-as-trace-of-program and formalize it into an operational semantics for our higher-order language. This semantics is given externally as an abstract machine, and internally, as a *monadic interpretation*.

The strength of our approach to circuit synthesis is to be able to reason on a regular program independently from the constraints of the circuit construction. The approach we follow is similar to what is done in Geometry of synthesis [9] for hardware synthesis, but since the back-end we aim at is way simpler, we can devise a very natural and compact monadic operational semantics.

Contribution. The main contribution of this paper is a monadic presentation of Landauer embeddings [16] in the context of higher-order programs. Its main strength is its parametricity: a program really represents a *family* of circuits, parametrized on the size of the input. Furthermore, we demonstrate a compositional monadic procedure for generating a reversible circuit out of a regular, purely functional program. The generated circuit is then provably computing the same thing as the original program. This can be used to internalize the generation of a reversible circuit out of a functional program. It has been implemented in Quipper [23] and used for building complex quantum oracles.

Related works. From the description of a conventional function it is always possible to design a reversible circuit computing the function out of its truth table or properties thereof and several methods have been designed to generate compact circuits (see e.g. [24, 18, 12, 25, 7, 35]). However, if these techniques allow one to write reversible functions with arbitrary truth tables [34], they do not usually scale well as the size of the input grows.

Synthesis of reversible circuits can be seen as a small branch of the vast area of hardware synthesis. In general, hardware synthesis can be structural (description of the structure of the circuit) or behavioral (description of algorithm to encode). Functional programming languages have been used for both. On the more structural side one finds Lava [6], BlueSpec [20], functional netlists [22], *etc*. On the behavioral side we have the Geometry of Synthesis [9], Esterel [4], ForSyDe [26], *etc*. Two more recent contributions sitting in between structural and behavioral approaches are worth mentioning. First, the imperative, *reversible* synthesis language SyRec [36], specialized for reversible circuits. Then, Thomsen's proposal [33], allowing to represent a circuit in a functional manner, highlighting the behavior of the circuit out of its structure.

On the logic side, the geometry of interaction [10] is a methodology that can be used to turn functional programs into reversible computation [1, 9, 32]: it is based on the idea of turning a typing derivation into a reversible automaton.

There have also been attempts to design reversible abstract machines and to compile regular programs into reversible computation. For example, a reversible version of the SEMCD machine has been designed [15]. More recently, the compiler REVS [21] aims at compiling conventional computation into reversible circuits.

Monadic semantics for representing circuits is something relatively common, specially among the DSL community: Lava [6], Quipper [11], Fe-Si [5], *etc*. Other approaches use more sophisticated constructions, with type systems based on arrows [14] in order to capture reversibility.

In the present work, the language is circuit-agnostic, and the interest of the method lies more in the fact that the monadic semantics to build reversible circuits is completely *implicit* and only added at circuit-generation time, following the approach in [31], rather than in the choice of the language. Compared to [14], our approach is also parametric in the sense that a program does not describe one fixed-size circuit but a family of circuits, parametrized by the size of the input.

Plan of the paper. Section 2 presents the definition of reversible circuits and how to perform computation with them. Section 3 describes a PCF-like lambda-calculus and proposes two operational semantics: one as a simple beta-reduction and one using an abstract machine and a partial evaluation procedure generating a circuit. Section 4 describes the call-by-value reduction strategy and explains how to internalize the abstract-machine within the language using a monad. Section 5 discusses the use of this technique in the context of the generation of quantum oracles, and discusses the question of optimizing the resulting circuits. Finally, Section 6 concludes and proposes some future investigations.

2 Reversible circuits

A reversible boolean circuit consists in a set of *open wires* and *elementary gates* attached onto the wires. Schematically, a reversible boolean circuit is of the form shown on the right. To each gate is associated a boolean operation, supposed to be reversible. In this circuit example, G is a one-bit operation (for example a not-gate, flipping a bit) while F is a two-bit operation. In each wire, a bit "flows" from left to right. All the bits go at the same pace. When a gate is met, the corresponding operation is applied on the wires attached to the gate. Since the gates are reversible, the overall circuit is reversible by making the bits flow backward.

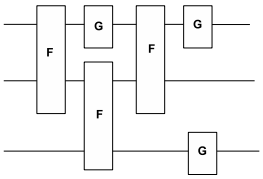

Choice of elementary gates. Many gates have been considered in the literature [24]. In this paper, we will consider multi-controlled-not gates. A not gate, represented by —⊕— , is flipping the value of the wire on which it is attached. The operator `not` stands for the bit-flip operation. Given a gate F acting on n wires, a controlled-F is a gate acting on $n+1$ wires. The control can be positive or negative, represented respectively as shown on the right. In both cases, the top wire is not modified. On the bottom wires, the gate F is applied if x is true for the positive control, and false for the negative control. Otherwise, no gate is applied: the values $\vec{y}$ flow unchanged through the gate. A positively-controlled not gate will be denoted CNOT.

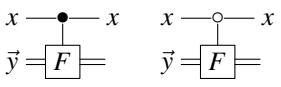

A reversible circuit runs a computation on some query: some input wires correspond to the query, and some output wires correspond to the answer. The auxiliary input wires that are not part of the query are initially fed with the boolean "false" (also written 0).

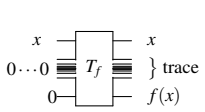

(a) Landauer embedding of f.

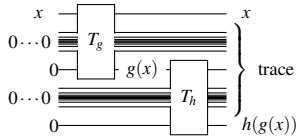

(b) Composing two Landauer embeddings.

Figure 1: Landauer embeddings.

Computing with reversible circuits. As described by Landauer [16] and Bennett [3], a conventional, classical algorithm that computes a boolean function $f : \texttt{bit}^n \to \texttt{bit}^m$ can be mechanically transformed into a reversible circuit sending the triplet $(x, \vec{0}, \vec{0})$ to $(x, \text{trace}, f(x))$, as in Figure 1a. Its input wires are not modified by the circuit, and the trace of all intermediate results are kept in garbage wires.

Because of their particular structure, two Landauer embeddings T_g and T_h can be composed to give a Landauer embedding of the composition $h \circ g$. Figure 1b shows the process: the wires of the output of T_g are fed to the input of T_h, and the output of the global circuit is the one of T_h. The garbage wires now contain all the ones of T_g and T_h.

Note that it is easy to build elementary Landauer embeddings for negation and conjunction: the former is a negatively-controlled not while the latter is a positively doubly-controlled not. Any boolean function can then be computed with Landauer embeddings.

3 Reversible circuits as trace of programs

In this section, we present an implementation of Landauer embeddings to the context of a higher-order functional programming language, and show how it can be understood through an abstract machine.

3.1 Simple formalization of reversible circuits

A reversible circuit has a very simple structure. As a linear sequence of elementary gates, it can be represented as a simple list of gates.

Definition 1. A *reversible gate* G is a term $\texttt{N}(i \cdot b_1^{j_1} \dots b_n^{j_n})$ where i, $j_1, \dots, j_n$ are natural numbers such that for all k, $i \neq j_k$, and where $b_1, \dots, b_n$ are booleans. If the list of $b_k^{j_k}$ is empty, we simply write $\texttt{N}(i)$ in place of $\texttt{N}(i \cdot)$. The *wires* of the gate $\texttt{N}(i \cdot b_1^{j_1} \dots b_n^{j_n})$ is the set of natural numbers $\{i, j_1, \dots, j_n\}$. The wire i is called *active* and the j_k's are called the *control wires*. Given a list C of gates, the union of the sets of wires of the elements of C is written $\mathit{Wires}(C)$. Finally, the boolean values True and False flowing in the wires are respectively represented with `tt` and `ff` throughout the paper.

Definition 2. A *reversible boolean circuit* is a triplet (I, C, O) where C is a list of reversible gates and where I and O are sets of wires. The list C is the *raw circuit*, I is the set of *inputs wires* and O the set of *outputs wires*. We also call $\mathit{Wires}(C) \setminus I$ the *auxiliary wires* and $\mathit{Wires}(C) \setminus O$ the *garbage wires*.

Executing a reversible circuit on a given tuple of booleans computes as follows.

Definition 3. Consider a circuit (I,C,O) and a family of bits $(x_i)_{i\in I}$. A *valuation* for the circuit is an indexed family $v = (v_j)_{j\in Wires(C)\cup I\cup O}$ of booleans. The *execution of a gate* $\mathtt{N}(i\cdot b_1^{j_1}\dots b_n^{j_n})$ *on the valuation* v is the valuation w such that for all $l\neq i$, $w_l = v_l$ and $w_i = v_i$ $\mathtt{xor}$ $\wedge_{k=1}^{n}(v_{j_k}$ $\mathtt{xor}$ b_k $\mathtt{xor}$ $\mathtt{tt})$ if $n\geq 1$ and $w_i = \mathtt{not}(v_i)$ otherwise. The execution of the circuit (I,C,O) with input $(x_i)_{i\in I}$ is the succession of the following operations: (1) Initialization of a valuation v such that for all $k\in I$, $v_k = x_k$, and for all the other values of k, v_k is false. (2) Execution of every gate in C on v, *in reverse order*. (3) The execution of the circuit returns the sub-family $(v_k)_{k\in O}$.

3.2 A PCF-like language with lists of booleans

In this section, we present the functional language $\mathbf{PCF}^{\mathbf{list}}$ that we use to describe the regular computations that we eventually want to perform with a reversible circuit. The language is simply-typed and it features booleans, pairs and lists.

$$\begin{aligned} M,N \;::=\;& x \mid \lambda x.M \mid MN \mid \langle M,N\rangle \mid \pi_1(M) \mid \pi_2(M) \mid \mathtt{skip} \mid M;N \mid \mathtt{tt} \mid \mathtt{ff} \mid \\ & \mathtt{if}\ M\ \mathtt{then}\ N\ \mathtt{else}\ P \mid \mathtt{and} \mid \mathtt{xor} \mid \mathtt{not} \mid \mathtt{inj}_1(M) \mid \mathtt{inj}_2(M) \mid \\ & \mathtt{match}\ P\ \mathtt{with}\ (x\mapsto M \mid y\mapsto N) \mid \mathtt{split}^A \mid Y(M) \mid \mathtt{Err}, \\ A,B \;::=\;& \mathtt{bit} \mid A\oplus B \mid A\times B \mid \mathtt{1} \mid A\to B \mid [A]. \end{aligned}$$

The language comes equipped with the typing rules of Table 1. There are several things to note. First, the construct `if-then-else` can only output *first-order types*. A first order type is a type from the grammar $A^0,B^0 ::= \mathtt{bit} \mid A^0\times B^0 \mid [A^0]$. Despite the fact that one can encode them with the test construct, for convenience we add the basic boolean combinators `not`, `xor` and `and`. There are no constructors for lists, but instead there is a coercion from $\mathtt{1}\oplus(A\times[A])$ to $[A]$; the term `split` turns a list-type into a additive type. There is a special-purpose term `Err` that will be used in particular in Section 3.4 as an error-spawning construct. The boolean values True and False are respectively represented with `tt` and `ff`. `skip` is the unit term and $M;N$ is used as the destructor of the unit. Finally, Y is a fixpoint operator. As we shall eventually work with a call-by-value reduction strategy, we only consider fixpoints defining functions.

Notation 4. We write `nil` for $\mathtt{inj}_1(\mathtt{skip})$ and $M :: N$ in place of $\mathtt{inj}_2(M\times N)$. We also write $[M_1,\dots M_n]$ for $M_1 :: \dots :: M_n :: \mathtt{nil}$. We also write general products $\langle M_1,\dots,M_n\rangle$ as $\langle M_1,\langle\dots M_n\dots\rangle\rangle$. Projections π_i for $i\leq n$ extends naturally to n-ary products. We write $\mathtt{letrec}\ f\,x = M\ \mathtt{in}\ N$ for the term $(\lambda f.N)(Y(\lambda f.\lambda x.M))$.

Remark 5. The typing rule of the `if-then-else` construct imposes a first-order condition on the branches of the test. This will be clarified in Remark 19. For now, let us just note that this constraint can be lifted with some syntactic sugar: if M and N are of type $A_1\to\dots\to A_n$, where A_n is first-order, then a "higher-order" test $\mathtt{if}\ P\ \mathtt{then}\ M\ \mathtt{else}\ N$ can be defined using the native first-order test by an η-expansion with the lambda-abstraction $\lambda x_1\dots x_n.\mathtt{if}\ P\ \mathtt{then}\ Mx_1\dots x_n\ \mathtt{else}\ Nx_1\dots x_n$.

$$\frac{}{\Delta, x : A \vdash x : A} \quad \frac{}{\Delta \vdash \mathtt{tt} : \mathtt{bit}} \quad \frac{}{\Delta \vdash \mathtt{ff} : \mathtt{bit}} \quad \frac{}{\Delta \vdash \mathtt{skip} : 1} \quad \frac{}{\Delta \vdash \mathtt{Err} : A}$$

$$\frac{}{\Delta \vdash \mathtt{not} : \mathtt{bit} \to \mathtt{bit}} \quad \frac{}{\Delta \vdash \mathtt{and} : \mathtt{bit} \times \mathtt{bit} \to \mathtt{bit}} \quad \frac{}{\Delta \vdash \mathtt{xor} : \mathtt{bit} \times \mathtt{bit} \to \mathtt{bit}}$$

$$\frac{}{\Delta \vdash \mathtt{split} : [A] \to 1 \oplus (A \times [A])}$$

$$\frac{\Delta, x : A \vdash M : B}{\Delta \vdash \lambda x.M : A \to B} \quad \frac{\Delta \vdash M : A_1 \times A_2}{\Delta \vdash \pi_i(M) : A_i} \quad \frac{\Delta \vdash M : A_i}{\Delta \vdash \mathtt{inj}_i(M) : A_1 \oplus A_2}$$

$$\frac{\Delta \vdash M : A \to B \quad \Delta \vdash N : A}{\Delta \vdash MN : B} \quad \frac{\Delta \vdash M : A \quad \Delta \vdash N : B}{\Delta \vdash \langle M, N \rangle : A \times B} \quad \frac{\Delta \vdash M : 1 \quad \Delta \vdash N : B}{\Delta \vdash M; N : B}$$

$$\frac{\Delta \vdash P : A \oplus B \quad \Delta, x : A \vdash M : C \quad \Delta, y : B \vdash N : C}{\Delta \vdash \mathtt{match}\ P\ \mathtt{with}\ (x^A \mapsto M | y^B \mapsto N) : C} \quad \frac{\Delta \vdash M : 1 \oplus (A \times [A])}{\Delta \vdash M : [A]}$$

$$\frac{\Delta \vdash M : A \to A}{\Delta \vdash Y(M) : A} \quad \frac{\Delta \vdash P : \mathtt{bit} \quad \Delta \vdash M : C \quad \Delta \vdash N : C \quad \text{the type } C \text{ is first-order}}{\Delta \vdash \mathtt{if}\ P\ \mathtt{then}\ M\ \mathtt{else}\ N : C}$$

Table 1: Typing rules of **PCF$^{\mathbf{list}}$**.

$$(\lambda x.M)N \to M[N/x] \qquad \pi_i \langle M_1, M_2 \rangle \to M_i \qquad \mathtt{skip}; M \to M$$

$$\mathtt{if}\ \mathtt{tt}\ \mathtt{then}\ M\ \mathtt{else}\ N \to M \qquad \mathtt{if}\ \mathtt{ff}\ \mathtt{then}\ M\ \mathtt{else}\ N \to N \qquad \mathtt{split}\ M \to M$$

$$\mathtt{match}\ \mathtt{inj}_i(P)\ \mathtt{with}\ (x_1 \mapsto M_1 | x_2 \mapsto M_2) \to M_i[P/x_i] \qquad Y(M) \to M(Y(M))$$

Table 2: Small-step semantics for **PCF$^{\mathbf{list}}$**: reduction rules, acting on subterms.

3.3 Small-step semantics

We equip the language **PCF$^{\mathbf{list}}$** with the smallest rewrite-system closed under subterm reduction, satisfying the rewrite rules of Table 2, and satisfying the obvious rules regarding `not`, `and` and `xor`: for example, `not tt` $\to$ `ff` and `not ff` $\to$ `tt`. Note that the term `Err` does not reduce. This is on purpose: it represents an error that one cannot catch with the type system; in particular it will be used in Section 3.4. The usual safety properties are satisfied, modulo the error-spawning term `Err`.

Definition 6. A value is a term V defined by the grammar $\lambda x.M \mid \langle U_1, U_2 \rangle \mid \mathtt{inj}_i(U) \mid c$, where c is a constant term: `skip`, `tt`, `ff`.

Theorem 7 (Safety)**.** *Type preservation and progress are verified: (1) If $\Delta \vdash M : A$, then for all N such that $M \to N$ we also have $\Delta \vdash N : A$. (2) If M is a closed term of type A then either M is a value, or M contains the term* `Err`*, or M reduces.* □

In summary, the language is well-behaved. It is also reasonably expressive, in the sense that most of the computations that one could want to perform on lists of bits can be described, as shown in Example 9.

Convention 8. When defining a large piece of code, we will be using a Haskell-like notation. So instead of defining a closed function as a lambda-term on a typing judgment, we shall be using the notation

```
function : type_of_the_function
function arg1 arg2 ... = body_of_the_function
```

Also, we shall use the convenient notation *let* $x = M$ *in* N for $(\lambda x.N)M$ and the notation *let* $\langle x,y\rangle = M$ *in* N for *let* $z = M$ *in let* $x = \pi_1(z)$ *let* $y = \pi_1(z)$ *in* N. Similarly, we allow multiple variables for recursive functions, and we use pattern-matching for lists and general products in the same manner.

Example 9 (List combinators)**.** The usual list combinators can be defined. Here we give the definition of `foldl:` `(A → B → A) → A → [B] → A`. The other ones (such as map, zip…) are written similarly.

```
foldl f a l = letrec g z l' = match (split l') with
                    nil   ↦ z
                  | ⟨h,t⟩ ↦ g (f z h) t
              in g a l
```

Example 10 (Ripple-carry adder)**.** One can easily encode a bit-adder: it takes a carry and two bits to add, and it replies with the answer and the carry to forward.

```
bit_adder : bit → bit → bit → (bit × bit)
bit_adder carry x y =
      let majority a b c = if (xor a b) then c else a in
      let z = xor (xor carry x) y in
      let carry' = majority carry x y in ⟨carry', z⟩
```

Encoding integers as lists of bits, low-bit first, one can use the bit-adder to write a complete adder in a ripple-carry manner, amenable to a simple folding. We use an implementation similar to the one done in [23].

```
adder_aux : (bit × [bit]) → (bit × bit) → (bit × [bit])
adder_aux ⟨w, cs⟩ ⟨a, b⟩ = let ⟨w', c'⟩ = bit_adder w a b in ⟨w', c'::cs⟩

adder : [bit] → [bit] → [bit]
adder x y = π₂ (foldl adder_aux ⟨ff, nil⟩ (zip y x))
```

3.4 Reversible circuits from operational semantics

We consider the language $\mathbf{PCF}^{\mathbf{list}}$ as a *specification language* for boolean reversible circuits in the following sense: A term of type $x_1 : \mathtt{bit}, \ldots, x_n : \mathtt{bit} \vdash M : \mathtt{bit}^m$ computes a boolean function $f_M : \mathtt{bit}^n \to \mathtt{bit}^m$.

In this section, we propose an operational semantics for the language $\mathbf{PCF}^{\mathbf{list}}$ generating Landauer embeddings, as described in Section 2. The circuit is produced during the execution of an abstract machine and partial evaluation of terms. Essentially, a term reduces as usual, except for the term constructs handling the type `bit`, for which we only record the operations to be performed. Formally, the definitions are as follows.

Definition 11. A *circuit-generating abstract machine* is a tuple consisting of (1) a typing judgment $p_1 : \mathtt{bit}, \ldots, p_{n+k} : \mathtt{bit} \vdash M : \mathtt{bit}^m$; (2) a partial circuit $RC := (\{1,\ldots,n\}, C)$ where C is a list of gates; (3) a one-to-one linking function mapping the free variables p_i of M to the wires $\mathit{Wires}(C) \cup \{1,\ldots,n\}$.

Intuitively, $\{1,\ldots,n\}$ is the set of input wires. The set of output wires is not yet computed: we only get it when M is a value. If G is a gate, we write $G :: (I,C)$ for the partial circuit $(I, G :: C)$. Given a judgment $p_1 : \mathtt{bit}, \ldots, p_n : \mathtt{bit} \vdash M : \mathtt{bit}^m$, the

$$
\begin{aligned}
(C[(\lambda x.M)N], RC, L) &\to_{am} (C[M[N/x]], RC, L)\\
(C[\pi_i\langle M_1, M_2\rangle], RC, L) &\to_{am} (C[M_i], RC, L)\\
(C[\texttt{skip};M], RC, L) &\to_{am} (C[M], RC, L)\\
(C[\texttt{split}\, M], RC, L) &\to_{am} (C[M], RC, L)\\
(C[\texttt{match inj}_i(P)\ \texttt{with}\ (x_1 \mapsto M_1 | x_2 \mapsto M_2)], RC, L) &\to_{am} (C[M_i[P/x_i]], RC, L)\\
(C[Y(M)], RC, L) &\to_{am} (C[M(Y(M))], RC, L)
\end{aligned}
$$

Table 3: Rewrite rules for circuit-generating abstract-machine: generic rules.

$$
\begin{aligned}
(C[\texttt{ff}], RC, L) \to_{am} (C[p_{i_0}], RC, L') \quad (C[\texttt{tt}], RC, L) &\to_{am} (C[p_{i_0}], (\texttt{N}(i_0)) :: RC, L')\\
(C[\texttt{not}\ p_i], RC, L) &\to_{am} (C[p_{i_0}], \texttt{N}(i_0 \cdot \texttt{ff}^i) :: RC, L')\\
(C[\texttt{and}\ p_i\ p_j], RC, L) &\to_{am} (C[p_{i_0}], \texttt{N}(i_0 \cdot \texttt{tt}^i\texttt{tt}^j) :: RC, L')\\
(C[\texttt{xor}\ p_i\ p_j], RC, L) &\to_{am} (C[p_{i_0}], \texttt{N}(i_0 \cdot (i, \texttt{tt})) :: \texttt{N}(i_0 \cdot \texttt{tt}^j) :: RC, L')\\
(C[\texttt{if}\ p_i\ \texttt{then}\ V\ \texttt{else}\ W], RC, L) &\to_{am}\\
&\begin{cases} (C[U], RC', L'') & V \text{ and } W \text{ of the same shape}\\ C[\texttt{Err}] & \text{otherwise}\end{cases}
\end{aligned}
$$

Table 4: Rewrite rules for circuit-generating abstract-machines: rules for booleans

empty machine is $(M, (\{1,\ldots,n\}, \{\}), \{p_i \mapsto i \quad | \quad i = 1\ldots n\})$ and is denoted with $\mathit{EmptyAM}(M)$. The size of the domain of a linking function L is written $\sharp(L)$.

By abuse of notation, we shall write abstract machine with terms, and not typing judgements. It is assumed that all terms are well-typed according to the definition.

Definition 12. Given a linking function L, a *first-order extension of* L consists of a term of shape $M ::= p_i \,|\, \langle M_1, \ldots M_n\rangle \,|\, [M_1, \ldots M_n]$, where the p_i's are in the domain of L. We say that two first-order extensions of L have *the same shape* provided that they are both products with the same size or lists with the same size such as their components have pairwise the same shape.

The set of circuit-generating abstract machines is equipped with a rewrite-system $(\to_{am})$ defined using a notion of *beta-context* $C[-]$, that is, a term with a hole, as follows.

$$
\begin{array}{l}
[-] \,|\, \lambda x.C[-] \,|\, (C[-])N \,|\, M(C[-]) \,|\, \langle C[-], N\rangle \,|\, \langle M, C[-]\rangle \,|\\
\pi_1(C[-]) \,|\, \pi_2(C[-]) \,|\, C[-];N \,|\, M;C[-] \,|\, \texttt{if}\ C[-]\ \texttt{then}\ N\ \texttt{else}\ P \,|\\
\texttt{if}\ M\ \texttt{then}\ C[-]\ \texttt{else}\ P \,|\, \texttt{if}\ M\ \texttt{then}\ N\ \texttt{else}\ C[-] \,|\, \texttt{inj}_1(C[-]) \,|\, \texttt{inj}_2(C[-]) \,|\\
\texttt{match}\ C[-]\ \texttt{with}\ (x \mapsto M | y \mapsto N) \,|\, \texttt{match}\ P\ \texttt{with}\ (x \mapsto C[-] | y \mapsto N) \,|\\
\texttt{match}\ P\ \texttt{with}\ (x \mapsto M | y \mapsto C[-]) \,|\, Y(C[-]).
\end{array}
$$

The constructor $[-]$ is the *hole* of the context. Given a context $C[-]$ and a term M, we define $C[M]$ as the variable-capturing substitution of the hole $[-]$ by M.

The rewrite rules can then be split in two sets. The first set concerns all the term constructs unrelated to the type `bit`. In these cases, the state of the abstract machine is not modified, only the term is rewritten. The rules, presented in Table 3, are the same as for the small-step semantics of Table 2: apart from the two rules concerning `if-then-else`, all the others are the same.

The second set of rules concerns the terms dealing with the type `bit`, and can be seen as partial-evaluation rules: we only record in the circuit the operations that would need to be done. The rules are shown in Table 4. The linking function L' is $L \cup \{p_{i_0} \mapsto i_0\}$, where i_0 is a new wire. The variable p_{i_0} is assumed to be fresh. For the case of the `if-then-else`, we assume V and W are first-order extensions of L with the same shape. The term U is a first-order extension of L with the same shape as V and W containing only (pairwise-distinct) free variables and mapping to new distinct garbage wires. L'' is L updated with this new data. Suppose that V contains the variables $v_1, \dots v_k$, that W contains the variables $w_1, \dots w_k$ and that U contains the variables $u_1, \dots u_k$. Then RC' is RC with the following additional series of gates: $\mathtt{N}(u_j \cdot \mathtt{tt}^{p_i}\mathtt{tt}^{v_i}))$ and $\mathtt{N}(u_j \cdot \mathtt{ff}^{p_i}\mathtt{tt}^{w_i})$.

Remark 13. Note that the set I is never modified by the rules

Safety properties hold for this new semantics, in the sense that the only error uncaught by the type system is the term `Err` that might be spawned.

Theorem 14 (Type preservation)**.** *If $p_1 : \mathtt{bit}, \dots, p_{\sharp(L)} : \mathtt{bit} \vdash M : \mathtt{bit}^m$, if (M, RC, L) is an abstract machine and if $(M, RC, L) \to_{am} (N, RC', L')$, then we have the judgement $p_1 : \mathtt{bit}, \dots, p_{\sharp(L')} : \mathtt{bit} \vdash N : \mathtt{bit}^m$.* □

Theorem 15 (Progress)**.** *Suppose that $p_1 : \mathtt{bit}, \dots, p_{\sharp(L)} : \mathtt{bit} \vdash M : \mathtt{bit}^m$ is valid and that (M, RC, L) is an abstract machine. Then either M is a value, or M contains `Err`, or (M, RC, L) reduces through ($\to_{am}$).* □

3.5 Simulations

The abstract machine M generates a circuit computing the same function as the small-step reduction of M in the following sense.

Definition 16. Let $(M, (I, C), L)$ be an abstract machine. We write $C(M, (I, C), L)$ for the circuit defined as $(I, C, \mathrm{Range}(L))$. Let $(v_k)_{k \in \mathrm{Range}(L)}$ be the execution of the circuit $C(M, (I, C), L)$ on the valuation $\vec{u} = (u_i)_{i \in I}$. We define $T(M, (I, C), L)(\vec{u})$ as the term M where each free variable x has been replaced with $v_{L(x)}$.

Intuitively, if (M, RC, L) is seen as a term where some boolean operations have been delayed in RC, then $T(M, RC, L)$ corresponds to the term resulting from the evaluation of the delayed operations.

Theorem 17. *Consider a judgment $x_1 : \mathtt{bit}, \dots, x_n : \mathtt{bit} \vdash M : \mathtt{bit}^m$ and suppose that $\mathit{EmptyAM}(M) \to_{am}{}^* (\langle p_{i_1}, \dots p_{i_k} \rangle, (I, C), L)$. Then $k = m$, and provided that $\vec{u} = (b_i)_{i \in I}$, the term $T(\langle p_{i_1}, \dots p_{i_k} \rangle, (I, C), L)(\vec{u})$ is equal to $\langle c_1, \dots c_m \rangle$ if and only if the term let $\langle x_1, \dots x_n \rangle = \langle b_1, \dots b_n \rangle$ in M reduces to $\langle c_1, \dots c_m \rangle$.*

The proof is done using an invariant on a single step of the rewriting of abstract machines, stated as follows.

Lemma 18. *Consider a judgment $x_1 : \mathtt{bit}, \ldots, x_n : \mathtt{bit} \vdash M : \mathtt{bit}^m$ and suppose that $(M,(I,C),L) \to_{am} (N,(I,C'),L')$. Let $\vec{u} = (u_i)_{i\in I}$ be a valuation. Then either the term $T(M,(I,C),L)(\vec{u})$ is equal to $T(N,(I,C'),L')(\vec{u})$ if the rewrite corresponds to the elimination of a boolean* `tt` *or* `ff`*, or $T(M,(I,C),L)(\vec{u}) \to T(N,(I,C'),L')(\vec{u})$, or N contains the error term* `Err`*.* □

Proof of Theorem 17. If $\mathit{EmptyAM}(M) \to_{am}{}^* (\langle p_{i_1}, \ldots p_{i_k}\rangle, (I,C), L)$, then there is a sequence of intermediate rewrite steps where none of the terms involved is the term `Err`. From Lemma 18, one concludes that for all valuations $\vec{u}$ on I, $T(\mathit{EmptyAM}(M))(\vec{u}) \to^* T(\langle p_{i_1}, \ldots p_{i_k}\rangle, (I,C), L)(\vec{u})$. Choosing $\vec{u} = (b_i)_{i\in I}$, $T(\mathit{EmptyAM}(M))(\vec{u})$ is the term M where each free variable p_{i_j} has been substituted with its corresponding boolean b_{i_j}. Similarly, $T(\langle p_{i_1}, \ldots p_{i_k}\rangle, (I,C), L)$ is equal to the value $\langle b_{i_1}, \ldots b_{i_k}\rangle$. We can conclude the proof by remarking that the term $\mathit{let}\ \langle x_1, \ldots x_n\rangle = \langle b_1, \ldots b_n\rangle\ \mathit{in}\ M$ reduces to M where each of the free variables p_{i_j} have been substituted with b_{i_j}, that is, the term $T(\mathit{EmptyAM}(M))(\vec{u})$. □

One would have also hoped to have a simulation result in the other direction, stating that if a (closed) term $M : \mathtt{bit}^m$ reduces to a tuple of booleans, then $\mathit{EmptyAM}(M)$ generates a circuit computing the same tuple. Unfortunately this is not the case, and the reason is the particular status of the type `bit` and the way the `if-then-else` behaves.

Remark 19. Let us re-visit the first-order constraint of the `if-then-else` discussed in Remark 5 in the light of this operational semantics. Here, this test behaves as a regular boolean operator acting on three arguments: they need to be all reduced to values before continuing. This test is "internal" to the circuit: both branches are evaluated during a run of the program. Because it is "internal", the type of the branches have to be "representable": thus the constraint on first-order. This test does not control the execution of the program: its characteristic only appears at circuit-evaluation time.

With this operational semantics, it is also interesting to note that there are two kinds of booleans: the "internal" type `bit`, and the "external" type defined e.g. as $\mathtt{bool} = \mathtt{1} \oplus \mathtt{1}$. If the former does not control the flow, the latter does with the `match` constructor. And unlike `if-then-else`, `match` does not have type constraints on its branches.

The term `Err` can be explained in the light of this discussion. Thanks to the condition on the shape of the output branches of the test, it is used to enforce the fact that `bit` cannot be coerced to a `bool`. Indeed, consider the term `if` b `then nil else` [`tt`]: using a `match` against the result of the test, it would allow one to use the bit b for controlling the shape of the rest of the circuit. As there is not such construct for reversible circuits, it therefore has to be forbidden: it is not possible to control the flow of execution of the program through the type `bit`. And the fact that a well-typed term can produce an error is simply saying that the type-system is not "strong enough" to capture such a problem. It is very much related to the fact that the `zip` operator on lists cannot be "safely" typed without dependent types.

4 Internalizing the abstract machine

Instead of defining an external operational semantics as we did in Section 3.4, one can internalize the definition of circuits in the language **PCF**$^{\textbf{list}}$. Given a program, provided that one chooses a reduction strategy, one can simulate the abstract-machine semantics inside **PCF**$^{\textbf{list}}$ using a generic *monadic lifting*, close to what was proposed in [31].

4.1 Monadic lifting

Before going ahead with the full abstract-machine semantics, we present the monadic lifting of **PCF**$^{\textbf{list}}$ for a monadic function-type. It is the transposition of Haskell's monads to our language **PCF**$^{\textbf{list}}$. The main characteristic of the reversible abstract-machine is to change the operational behavior of the type `bit`: the terms `tt`, `ff`, the inline `bit`-combinators and the term construct `if-then-else` do not reduce as regular lambda-terms. Instead, they trigger a side-effect, which can be simulated within a monad.

Definition 20. A *monad* is a function-type $\mathscr{M}(-)$ together with two terms $\texttt{return}^A_{\mathscr{M}} : A \to \mathscr{M}(A)$ and $\texttt{app}^{A,B}_{\mathscr{M}} : \mathscr{M}(A) \to (A \to \mathscr{M}(B)) \to \mathscr{M}(B)$. A *reversible-circuit monad* is a monad together with a type `wire` and the terms $\texttt{mtt}_{\mathscr{M}}, \texttt{mff}_{\mathscr{M}} : \mathscr{M}(\texttt{wire})$, $\texttt{mif}^A_{\mathscr{M}} : \texttt{wire} \to \mathscr{M}(A) \to \mathscr{M}(A) \to \mathscr{M}(A)$, and $\texttt{mnot}^A_{\mathscr{M}} : \mathscr{M}(\texttt{wire} \to \mathscr{M}(\texttt{wire}))$, and finally $\texttt{mand}^A_{\mathscr{M}}, \texttt{mxor}^A_{\mathscr{M}} : \mathscr{M}(\texttt{wire} \times \texttt{wire} \to \mathscr{M}(\texttt{wire}))$.

Definition 21. Given a reversible-circuit monad $\mathscr{M}$, we inductively define the *$\mathscr{M}$-monadic lifting of a type A*, written $\mathit{Lift}_{\mathscr{M}}(A)$. We omit the index $\mathscr{M}$ for legibility.

$$\begin{aligned} \mathit{Lift}(\texttt{bit}) &= \texttt{wire}, & \mathit{Lift}(\texttt{1}) &= \texttt{1}, \\ \mathit{Lift}(A \to B) &= \mathit{Lift}(A) \to \mathscr{M}(\mathit{Lift}(B)), & \mathit{Lift}(A \times B) &= \mathit{Lift}(A) \times \mathit{Lift}(B), \\ \mathit{Lift}(A \oplus B) &= \mathit{Lift}(A) \oplus \mathit{Lift}(B), & \mathit{Lift}([A]) &= [\mathit{Lift}(A)]. \end{aligned}$$

The *$\mathscr{M}$-monadic lifting of a term M*, denoted with $\mathit{Lift}_{\mathscr{M}}(M)$, is defined as follows. First, we set $\mathit{Lift}(\texttt{tt}) = \texttt{mtt}$, $\mathit{Lift}(\texttt{ff}) = \texttt{mff}$, $\mathit{Lift}(\texttt{and}) = \texttt{mand}$, $\mathit{Lift}(\texttt{xor}) = \texttt{mxor}$ and $\mathit{Lift}(\texttt{not}) = \texttt{mnot}$. Then

$$\begin{aligned} \mathit{Lift}(x) &= \texttt{return}\ x, \quad \mathit{Lift}(\texttt{skip}) = \texttt{return skip}, \\ \mathit{Lift}(\lambda x.M) &= \texttt{return}\ \lambda x.\mathit{Lift}(M), \quad \mathit{Lift}(\texttt{split}) = \texttt{return}\ \lambda \texttt{x.return}\ (\texttt{split x}), \\ \mathit{Lift}(MN) &= \texttt{app}\ \mathit{Lift}(M)\ \lambda x.\texttt{app}\ \mathit{Lift}(N)\ \lambda y.xy, \\ \mathit{Lift}(\langle M,N\rangle) &= \texttt{app}\ \mathit{Lift}(M)\ \lambda x.\texttt{app}\ \mathit{Lift}(N)\ \lambda y.\texttt{return}\ \langle x,y\rangle, \\ \mathit{Lift}(\pi_i(M)) &= \texttt{app}\ \mathit{Lift}(M)\ \lambda x.\texttt{return}\ \pi_i(x), \\ \mathit{Lift}(\texttt{inj}_i(M)) &= \texttt{app}\ \mathit{Lift}(M)\ \lambda x.\texttt{return}\ \texttt{inj}_i(x), \\ \mathit{Lift}(M;N) &= \texttt{app}\ \mathit{Lift}(M)\ \lambda x.\texttt{app}\ \mathit{Lift}(N)\ \lambda y.\texttt{return}\ x;y, \\ \mathit{Lift}(\texttt{match}\ P\ \texttt{with}\ (z_1 \mapsto M | z_2 \mapsto N)) &= \\ & \texttt{app}\ \mathit{Lift}(P)\ \lambda x.\texttt{match}\ x\ \texttt{with}\ (z_1 \mapsto \mathit{Lift}(M) | z_2 \mapsto \mathit{Lift}(N)), \\ & \mathit{Lift}(Y(M)) = \\ & \texttt{app}\ \mathit{Lift}(M)\ \lambda f.\texttt{return}\ (Y(\lambda y.\lambda z.\texttt{app}\,(y\,\texttt{skip})\,f))\texttt{skip} \\ \mathit{Lift}(\texttt{if}\ P\ \texttt{then}\ M\ \texttt{else}\ N) &= \texttt{app}\ \mathit{Lift}(P)\ \lambda x.((\texttt{mif}\ x)\ \mathit{Lift}(M))\ \mathit{Lift}(N) \end{aligned}$$

Remark 22. Note that in this definition of the lifting, we followed a call-by-value approach: the argument $N : \mathit{Lift}_{\mathscr{M}}(A)$ of a function $M : A \to \mathit{Lift}_{\mathscr{M}}(B)$ is first reduced to a value before being fed to the function. This will be discussed in Section 4.3.

The fact that a monad is equipped with `mtt`, `mff`, `mxor`, `mand`, `mnot` and `mif` is not a guarantee that the lifting will behave as expected. One has to choose the right monad for it. It is the topic of Section 4.2. However, in general this monadic-lifting operation preserves types (proof by induction on the typing derivation).

Theorem 23. *Provided that* $x_1 : A_1, \ldots, x_n : A_n \vdash M : B$ *is valid, so is the judgment* $x_1 : \mathit{Lift}_{\mathscr{M}}(A_1), \ldots, x_n : \mathit{Lift}_{\mathscr{M}}(A_n) \vdash \mathit{Lift}_{\mathscr{M}}(M) : \mathscr{M}(\mathit{Lift}_{\mathscr{M}}(B))$. □

4.2 Reversible circuits from monadic lifting

All the structure of the abstract machine can be encoded in the language $\mathbf{PCF}^{\mathbf{list}}$. A wire is a natural number. A simple way to represent them is with the type $\mathtt{wire} := [\mathtt{1}]$. The number 0 is the empty list while the successor of n is $(\mathtt{skip} :: n)$. A gate is then $\mathtt{gate} := \mathtt{wire} \times [\mathtt{wire} \times \mathtt{bit}]$. A raw circuit is $[\mathtt{gate}]$.

We now come to the abstract machine. In the formalization of Section 3.4, we were using a state with a circuit and a linking function. In this internal representation, the linking function is not needed anymore: the computation directly acts on wires. However, the piece of information that is still needed is the next fresh value. The state is encapsulated in $\mathtt{state} := [\mathtt{gate}] \times \mathtt{wire}$. Finally, given a type A, we write $\mathtt{circ}(A)$ for the type $\mathtt{state} \to (\mathtt{state} \times A)$: this is a computation generating a reversible circuit.

The type operator $\mathtt{circ}(-)$ can be equipped with the structure of a reversible-circuit monad, as follows. First, it is obviously a state-monad, making the two first constructs automatic:

$$\mathtt{return} := \lambda x.\lambda s.(s,x) \text{ and } \mathtt{app} := \lambda x f.\lambda s.\mathit{let}\ \langle s', a\rangle\ =\ x s\ \mathit{in}\ (f\,a)s'.$$

The others are largely relying on the fact that $\mathbf{PCF}^{\mathbf{list}}$ is expressive enough to emulate what was done in Section 3.4. Provided that $\mathbf{S}$ stands for the successor function, we can `mff` as the lambda-term $\lambda s.\mathit{let}\ \langle c, w\rangle\ = s\ \mathit{in}\ \langle\langle c, \mathbf{S}w\rangle, w\rangle$ and `mtt` as the lambda-term $\lambda s.\mathit{let}\ \langle c, w\rangle\ = s\ \mathit{in}\ \langle\langle\langle w, \mathtt{nil}\rangle :: c, \mathbf{S}w\rangle, w\rangle$. Note how the definition reflects the reduction rules corresponding to `tt` and `ff` in Table 4: in the case of `ff`, the returned wire is the next fresh one, and the state is updated by increasing the "next-fresh" value by one unit. In the case of `tt`, on top of this we add a not-gate to the list of gates in order to flip the value of the returned wire. The definitions of `mnot`, `mand` and `mxor` are similar. For `mif`, one capitalizes on the fact that we know the structure of the branches of the test, as they are of first-order types. One can then define a zip-operator $A^0 \times A^0 \to A^0$ for each first-order type A^0.

4.3 Call-by-value reduction strategy

As was mentioned in Remark 22, the monadic lifting intuitively follows a call-by-value approach. It can be formalized by developing a call-by-value reduction strategy

$$(E[(\lambda x.M)V],RC,L) \to_{cbv} (E[M[V/x]],RC,L)$$
$$(E[\pi_i\langle V_1,V_2\rangle],RC,L) \to_{cbv} (E[V_i],RC,L)$$
$$(E[\mathtt{skip};M],RC,L) \to_{cbv} (E[M],RC,L)$$
$$(E[\mathtt{split}\,V],RC,L) \to_{cbv} (E[V],RC,L)$$
$$(E[\mathtt{match}\ \mathtt{inj}_i(V)\ \mathtt{with}\ (x_1 \mapsto M_1 | x_2 \mapsto M_2)],RC,L) \to_{cbv} (E[M_i[V/x_i]],RC,L)$$
$$(E[Y(\lambda x.M)],RC,L) \to_{cbv} (E[M[Y(\lambda x.M)/x]],RC,L)$$

Table 5: Call-by-value for circuit-generating abstract-machine: generic rules.

$$(E[\mathtt{ff}],RC,L) \to_{cbv} (E[p_{i_0}],RC,L') \quad (E[\mathtt{tt}],RC,L) \to_{cbv} (E[p_{i_0}],(\mathtt{N}(i_0)) :: RC,L')$$
$$(E[\mathtt{not}\ p_i],RC,L) \to_{cbv} (E[p_{i_0}],\mathtt{N}(i_0 \cdot \mathtt{ff}^i) :: RC,L')$$
$$(E[\mathtt{and}\ p_i\ p_j],RC,L) \to_{cbv} (E[p_{i_0}],\mathtt{N}(i_0 \cdot \mathtt{tt}^i\mathtt{tt}^j) :: RC,L')$$
$$(E[\mathtt{xor}\ p_i\ p_j],RC,L) \to_{cbv} (E[p_{i_0}],\mathtt{N}(i_0 \cdot (i,\mathtt{tt})) :: \mathtt{N}(i_0 \cdot \mathtt{tt}^j) :: RC,L')$$
$$(E[\mathtt{if}\ p_i\ \mathtt{then}\ V\ \mathtt{else}\ W],RC,L) \to_{cbv} \begin{cases} (E[U],RC',L'') & V \text{ and } W \text{ of the same shape} \\ \mathtt{Err} & \text{otherwise} \end{cases}$$

Table 6: Call-by-value for circuit-generating abstract-machines: rules for booleans

for circuit-abstract machines. Such a definition follows the one of the reduction proposed in Section 3.4: we first design a notion of *call-by-value evaluation context* $E[-]$ characterizing the call-by-value redex that can be reduced.

Definition 24. A *call-by-value context* is a beta-context with the following grammar

$$\begin{array}{rcl} E[-] & ::= & [-] \,|\, (E[-])N \,|\, V(E[-]) \,|\, \langle E[-],N\rangle \,|\, \langle V,E[-]\rangle \,| \\ & & \pi_1(E[-]) \,|\, \pi_2(E[-]) \,|\, E[-];N \,|\, V;E[-] \,|\, \mathtt{if}\ E[-]\ \mathtt{then}\ N\ \mathtt{else}\ P \,| \\ & & \mathtt{if}\ V\ \mathtt{then}\ E[-]\ \mathtt{else}\ P \,|\, \mathtt{if}\ V\ \mathtt{then}\ W\ \mathtt{else}\ E[-] \,|\, \mathtt{inj}_1(E[-]) \,| \\ & & \mathtt{inj}_2(E[-]) \,|\, \mathtt{match}\ E[-]\ \mathtt{with}\ (x \mapsto M | y \mapsto N) \,|\, Y(E[-]). \end{array}$$

We define the *call-by-value reduction strategy* on circuit-generating abstract machines as shown in Tables 5 and 6.

Remark 25. Essentially, the generic rules of Table 3 are turned into their call-by-value version in the standard way. For example, we require that $(E[(\lambda x.M)V],RC,L) \to_{cbv} (E[M[V/x]],RC,L)$ happens only when V is a value. Similarly, the rules of Table 4 are reflected in Table 6, replacing $C[-]$ with $E[-]$.

Remark however that the reduction strategy does not exactly match the rewrite system $(\to_{am})$ in one special case:the rewrite rule concerning the fixpoint. We chose to modify it in order for the fixpoint to behave in a call-by-value manner: we embedded two steps of $(\to_{am})$ into one step of the strategy. One can then show that if $(M,RC,L) \to_{cbv} (M',RC',L')$ then $(M,RC,L) \to_{am}{}^* (M',RC',L')$.

In the light of this reduction strategy and of the monadic lifting of the previous section, one can now formalize what was mentioned in Remark 22. First, one can turn an abstract machine into a lifted term.

Definition 26. Let $x_1 : \mathtt{bit}, \ldots, x_{n+k} : \mathtt{bit} \vdash M : B$ and let $(M,(C,I),L)$ be an abstract machine where $I = \{1 \ldots n\}$. Then we define $\mathit{Lift}(M,(C,I),L)$ as the term

$$\Big(\mathit{Lift}(M)[\overline{L(x_{n+1})}/x_{n+1} \ldots \overline{L(x_{n+k})}/x_{n+k}]\Big)\ \langle \overline{C}, \mathbf{S}\,\overline{\max(\mathrm{Range}(L))}\rangle,$$

where $\overline{C}$ is the representation of C as a term of type $[\mathtt{gate}] \times \mathtt{wire}$, and where $\overline{n}$ with n an integer is the representation of n as a term of type $[1]$.

Then, provided that $\simeq_\beta$ stands for the reflexive, symmetric and transitive closure of the beta-reduction on terms and choosing M and $(M,(C,I),L)$ as in Definition 26:

Theorem 27. *Suppose that* $(M,RC,L) \to_{cbv} (M',RC',L')$. *Then* $\mathit{Lift}(M,RC,L)$ *is beta-equivalent to* $\mathit{Lift}(M',RC',L')$. □

Provided that the beta-reduction is confluent, this essentially says that the abstract-machine semantics can be simulated with the monadic lifting.

Corollary 28. *If* $\vdash M : \mathtt{bit}^m$ *and* $\mathit{EmptyAM}(M) \to_{cbv} (\langle x_1, \ldots x_m\rangle, (C,I), L)$, *then the term* $\pi_1(\mathit{Lift}(M))$ *is beta-equivalent to* $\overline{C}$, *where* $\overline{C}$ *is the representation of* C *as a term, as described in Definition 26.* □

5 Synthesis of quantum oracles

A rapid explanation is needed here: In quantum computation, one does not deal with classical bits but with the so-called *quantum bits*. At the logical level, a quantum algorithm consists of one or several *quantum circuits*, that is, reversible circuits with quantum bits flowing in the wires.

Quantum algorithms are used to solve classical problems. For example: factoring an integer [28], finding an element in a unordered list [19], finding the solution of a system of linear equations [13], finding a triangle in a graph [17], *etc*. In all of these algorithms, the description of the problem is a completely non-reversible function $f : \mathtt{bit}^n \to \mathtt{bit}^m$ and it has to be encoded as a reversible circuit computing the function $\bar{f} : \mathtt{bit}^n \times \mathtt{bit}^m \to \mathtt{bit}^n \times \mathtt{bit}^m$ sending (x,y) to $(x, y \,\mathtt{xor}\, f(x))$, possibly with some auxiliary wires set back to 0.

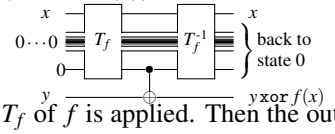

A canonical way to produce such a circuit is with a *Bennett embedding*. The procedure is shown on the right. First the Landauer embedding T_f of f is applied. Then the output of the circuit is `xor`'d onto the y input wires, and finally the inverse of T_f is applied. In particular, all the auxiliary wires are back to the value 0 at the end of the computation.

The method we propose in this paper offers a procedure for generating the main ingredient of this construction: the Landauer embedding. One just has to encode the problem in the language **PCF**$^{\mathbf{list}}$ (or extension thereof), possibly test and verify the program, and generate a corresponding reversible circuit through the monadic lifting. Theorems 17 and 28 guarantee that the monadic lifting of the program will give a circuit computing the same function as the original program.

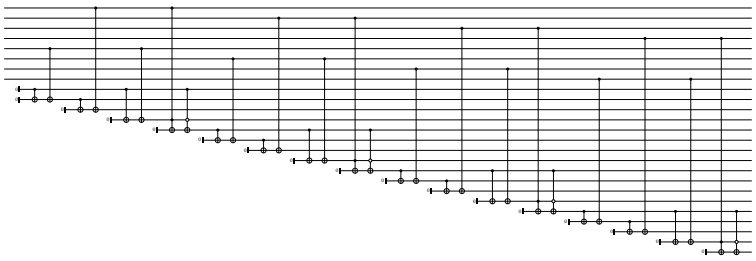

Figure 2: Reversible adder for 4-bit integers.

This algorithm was implemented within the language Quipper, and used for non-trivial oracles [23, 11]. Note that Quipper is not the only possible back-end for this generic monadic lifting: nothing forbids us to write a back-end in, say, Lava [6].

Example 29. The first example of code we saw (Example 10) computes an adder. One can run this code to generate a reversible adder: Figure 2 shows the circuit generated when fed with 4-bits integers. One can see 4 blocks of pairs of similar shapes.

Example 30. In the oracle for the QLSA algorithm [13, 27], one has to solve a system of differential equations coming from some physics problem using finite elements method. The bottom line is that it involves analytic functions such as sine and atan2.

Using fixed-point real numbers on 64 bits, we wrote a sine function using a Taylor expansion approximation. In total, we get a reversible circuit of 7,344,140 multi-controlled gates (with positive and negative controls). The function atan2 was defined using the CORDIC method. The generated circuit contains 34,599,531 multi-controlled gates. These two functions can be found in Quipper's distribution [23].

5.1 Efficiency of the monadic lifting

The monadic lifting proposed in this paper generates circuits that are efficient in the sense that the size of a generated circuit is linear in the number of steps it takes to evaluate the corresponding program. This means that any program running in polynomial time upon the size of its input generates a polynomial-sized circuit. Without any modification or optimization whatsoever, the technique is therefore able to generate an "efficient" circuit for an arbitrary, conventional algorithm. This is how the circuit for the function sine cited in Example 30 was generated: first, a conventional implementation was written and tested. When ready the lifting was performed, generating a circuit.

5.2 Towards a complete compiler

Compared to other reversible compilers [21], the approach taken in this paper considers the construction of the circuit as a process that can be completely automatized: the stance is that it should be possible to take a classical, functional program with conventional inductive datatypes and let the compiler turn it into a reversible circuit without

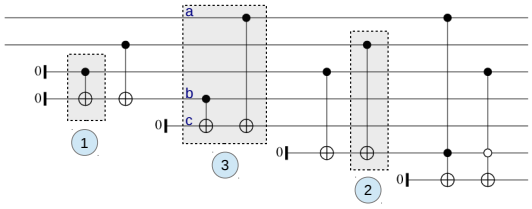

Figure 3: Adder of 1-bit integers, with potential optimizations highlighted

having to interfer (or only marginally). We do not claim to have a final answer: we only aim at proposing a research path towards such a goal.

A first step towards a more complete compiler for $\mathbf{PCF}^{\mathbf{list}}$ would involve optimization passes on the generated circuits. Indeed, as can be inferred from a quick analysis of Figure 2, if monadic lifting generates efficient circuits it does not produces particularly lean circuits.

5.3 Towards optimization of generated circuits

There is a rich literature on optimization of reversible circuits [24, 29, 18], some also considering positive and negative controls [30]. All of these works are relevant for reducing the size of the circuits we get.

The purpose of this section is not to aggressively optimize the circuits we get from the monadic lifting, but instead to reason on the particular shapes we obtain from this monadic semantics. If the reduction of a lambda-term into a reversible circuit is so verbose, it is partly due to the fact that garbage wires are created for every single intermediate result. We aim at pointing out the few optimization rules stemming from the circuit generation and reflecting these low-level optimizations to high-level program transformations.

The reversible adder of Figure 2 is very verbose. By applying the simple optimization schemes presented in this section, one gets the smaller circuit of Figure 5. One clearly sees the carry-ripple structure, and it is in fact very close to known reversible ripple-carry adders (see e.g. [8]). In the following discussion, we hint at how program transformations could be applied in order to get a circuit of compactness similar to the one obtained from the low-level circuit rewrites.

Algebraic optimizations. Let us consider the example of the 1-bit adder of Figure 3, from the code of Example 10. Three simple potential optimizations are highlighted.

In general, these optimizations require to have a knowledge of the value of the bits flowing in the wires (e.g. Dashed Box 2). Since there are input wires, this information needs to be kept in algebraic form, as a function of the input wires. Of course, for non-trivial circuits this means actually computing the circuit.

However, because of the shape of the generated circuit, instead of a complete algebraic form, for the purpose of circuit simplification it is often enough to keep only partial algebraic information about the wires. To each piece of wire, we essentially keep a limited knowledge of the past operations.

Dashed Box 1. Gates acting on wires of known constant value. The gate will never fire as the control will always be negative. The gate can be removed.

From the perspective of the code the circuit comes from, this situation typically occurs when constant booleans are manipulated, for example with the term `xor ff ff`.

Dashed Box 2. Copy of one wire to another one. Provided that the controlled wire is never controlled later on, one can remove the gate and move all the controls and gates of the bottom wire to the top wire.

The fact that the wire is not used later on means that the particular intermediate result is never used again: From the point of view of the program it means that this particular result is only used *linearly*. The typical case where this occurs is in a term such as $\texttt{and}\,(\texttt{not}\,(\texttt{and}\,x\,y))\,z$. A garbage wire is created to hold the result of the `not`, but this is not needed as this intermediate result is not going to be reused. Instead one can directly apply a not-gate on the result of $\texttt{and}\,x\,y$.

Dashed Box 3. Here, $c = a\,\texttt{xor}\,b$. The two CNOTs can be replaced with only one connecting wires a and b, and one could have removed wire c altogether. Again, some care must be taken: the new active wire should not be controlled later on (ruling out wire a), and the controls and actions of wire c should be moved to the new active wire.

This situation can also be understood as a linearity constraint on the program side.

Reduction strategies and garbage wires. The call-by-value reduction strategy we follow sometimes computes unused intermediate results, therefore generating gates acting on unused wires. One can safely discard such gates.

Note that the abstract machine is agnostic to the choice of reduction strategy. In general, depending on the chosen reduction path, the generated circuits do not have the same size and shape. Consider the term $x : \texttt{bit} \vdash (\lambda y.\texttt{and}\ y\ y)(\texttt{not}\ x) : \texttt{bit}$. A call-by-value reduction strategy first evaluate the argument, and then feeds the output value to the `and` operator. Since having the same variable y means that the two controlling wires collapsed, we get the circuit presented on the right: $\texttt{not}\ x$ is the output of the circuit, which is the meaning of the lambda-term.

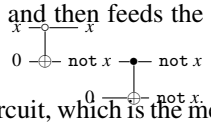

With a call-by-name strategy, the term instead reduces to the following circuit. The last wire is the output wire, the other wires are garbage wires. Of course, we get the same result: $(\texttt{not}\ x)\ \texttt{and}\ (\texttt{not}\ x)$ and $\texttt{not}\ x$ are indeed the same boolean value. However, note that the circuit is different than the one generated by a call-by-value strategy. In general, call-by-name tends to generate larger circuits as arguments are duplicated and evaluated several times. The case where it is not true is when the argument of a function is not used: in call-by-value, the argument would generate a piece of circuit, whereas in call-by-name, since it would not be evaluated, the argument would leave no trace on the circuit.

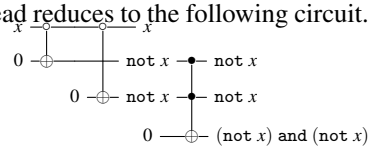

Optimizations by shuffling. A less obvious circuit modification is to send CNOT gates as far as possible to the right, by swapping the order of gates. This is again a side-effect of the particular shape of the generated circuit.

If it does not decrease the size the circuit, it is able to reveal places where algebraic optimizations can be performed. For example, consider the circuit in Figure 4a.

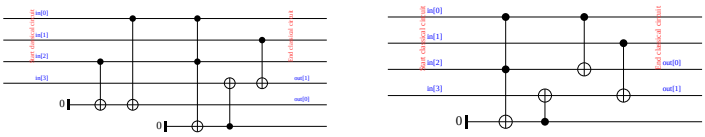

(a) Circuit waiting for shuffling. (b) After the shuffling.

Figure 4: Circuit optimized with and without shuffle.

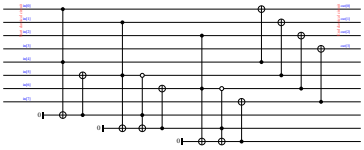

Figure 5: Reversible adder for 4-bit integers, optimized.

The two first CNOTs can be moved to the far right-end, becoming a hidden CNOT (as Dashed Box 2 of Figure 3): they are merged into one single CNOT and the first auxiliary wire is removed. We get the circuit in Figure 4b.

The corresponding program transformation modifies the term

$$\mathit{let}\ z_1 \ =\ \texttt{xor}\, x\, y\ \mathit{in}\ \mathit{let}\ z_2 \ =\ f(x,y)\ \mathit{in}\ g(z_1,z_2)$$

to

$$\mathit{let}\ z_2 \ =\ f(x,y)\ \mathit{in}\ \mathit{let}\ z_1 \ =\ \texttt{xor}\, x\, y\ \mathit{in}\ g(z_1,z_2).$$

As g does not use x nor y, one of the algebraic optimization might apply.

Example 31. By applying these optimization schemes on the reversible adder of Figure 2, one gets the circuit presented in Figure 5. One can now clearly see the carry-ripple structure, and it is in fact very close to known reversible ripple-carry adders (see e.g. [8]). These optimizations were implemented in Quipper: applied on larger circuits such as the ones of Example 30, we get in general a size reduction by a factor of 10.

6 Conclusion and future work

In this paper, we presented a simple and scalable mechanism to turn a higher-order program acting on booleans into into a family of reversible circuits using a monadic semantics. The main feature of this encoding is that an automatically-generated circuit is guaranteed to perform the same computation as the original program. The classical description we used is a small PCF-like language, but it is clear from the presentation that another choice of language can be made. In particular, an interesting question is whether it is possible to use a language with a stronger type system for proving properties of the encoded functions.

A second avenue of research is the question of the parallelization of the generated circuits. The circuits we produce are so far completely linear. Following the approach

in [9], using parallel higher-order language might allow one to get parallel reversible circuits, therefore generating circuits with smaller depths.

Finally, the last avenue for research is the design of generic compiler with a dedicated code optimizations. Indeed, an analysis of the specific optimizations described in Section 5.3 suggests that these could be designed at the level of code, therefore automatically generating leaner circuits up front. This opens the door to the design of specific type systems and code manipulations in a future full compiler. and back-end, specific circuit optimizations.

References

[1] S. Abramsky. A structural approach to reversible computation. *Theor. Comput. Sci.*, 347(3):441–464, 2005.

[2] A. Ambainis, A. M. Childs, *et al.* Any AND-OR formula of size n can be evaluated in time $n^{\frac{1}{2}+o(1)}$ on a quantum computer. *SIAM J. Comput.*, 39:2513–2530, 2010.

[3] C. H. Bennett. Logical reversibility of computation. *IBM J. Res. Dev.*, 17:525–532, 1973.

[4] G. Berry. The foundations of esterel. In *Proof, Language, and Interaction, Essays in Honour of Robin Milner*. MIT Press, 2000.

[5] T. Braibant and A. Chlipala. Formal verification of hardware synthesis. In *Computer Aided Verification*, volume 8044 of *LNCS*, pp. 213–228, 2013.

[6] K. Claessen. *Embedded Languages for Describing and Verifying Hardware*. PhD thesis, Chalmers University of Technology and Göteborg University, 2001.

[7] K. Fazel, M.A. Thornton, and J. E. Rice. ESOP-based Toffoli gate cascade generation. In *Proc. PacRim*, pp. 206–209, 2007.

[8] R. P. Feynman. Quantum mechanical computers. *Optics News*, 11:11–20, 1985.

[9] D. R. Ghica. Geometry of synthesis. In *Proc. POPL*, pp. 363–375, 2007.

[10] J.-Y. Girard. Towards a geometry of interaction. *Contemp. Math.*, 92:69–108, 1989.

[11] A. S. Green, P. L. Lumsdaine *et al.*, Quipper: a scalable quantum programming language. In *Proc. PLDI*, pp. 333–342, 2013.

[12] P. Gupta, A. Agrawal, and N. K. Jha. An algorithm for synthesis of reversible logic circuits. *IEEE Trans. on CAD of Int. Circ. and Sys.*, 25(11):2317–2330, 2006.

[13] A. W. Harrow, A. Hassidim, and S. Lloyd. Quantum algorithm for linear systems of equations. *Phys. Rev. Lett.*, 103(15):150502, 2009.

[14] Roshan P. James and A. Sabry. Information effects. In *Proc. POPL*, pp. 73–84, 2012.

[15] W. Kluge. A reversible SE(M)CD machine. In *Proc. IFL*, LNCS 1868, pp. 95–113, 1999.

[16] R. Laundauer. Irreversibility and heat generation in the computing process. *IBM J. Res. Dev.*, 5:261–269, 1961.

[17] F. Magniez, M. Santha, and M. Szegedy. Quantum algorithms for the triangle problem. *SIAM J. Comput.*, 37(2):413–424, 2007.

[18] D. Maslov, G. W. Dueck, and D. M. Miller. Fredkin/Toffoli templates for reversible logic synthesis. In *Proceedings of ICCAD*, pp. 256–261, 2003.

[19] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge Univ. Press, 2002.

[20] R. S. Nikhil. Bluespec: A general-purpose approach to high-level synthesis based on parallel atomic transactions. In *High-Level Synthesis*, pp. 129–146. Springer, 2008.

[21] A. Parent, M. Roetteler, and K. M. Svore. Reversible circuit compilation with space constraints. arXiv:1510.00377, 2015.

[22] S. Park, J. Kim, and H. Im. Functional netlists. In *Proc. ICFP*, pp. 353–366, 2008.

[23] Quipper. `http://www.mathstat.dal.ca/~selinger/quipper/`.

[24] M. Saeedi and I. L. Markov. Synthesis and optimization of reversible circuits – a survey. *ACM Comput. Surv.*, 45(2):21:1–21:34, March 2013.

[25] Y. Sanaee, M. Saeedi, and M. S. Zamani. Shared-pprm: A memory-efficient representation for boolean reversible functions. In *Proc. of ISVLSI*, pp. 471–474, 2008.

[26] I. Sander. *System Modeling and Design Refinement in ForSyDe*. PhD thesis, Royal Institute of Technology, Stockholm, Sweden, 2003.

[27] A. Scherer, B. Valiron *et al.* Resource analysis of the quantum linear system algorithm. arXiv:1505.06552, 2015.

[28] P Shor. Algorithms for quantum computation: discrete logarithm and factoring. In *Proc. FOCS*, 1994.

[29] M. Soeken, S. Frehse *et al.* RevKit: An open source toolkit for the design of reversible circuits. In *Proc. RC*, vol. 7165 of *LNCS*, pp. 64–76, 2012.

[30] M. Soeken and M. K. Thomsen. White dots do matter: Rewriting reversible logic circuits. In *Proc. RC*, vol. 7948 of *LNCS*, pp. 193–208, 2013.

[31] N. Swamy, N. Guts *et al.* Lightweight monadic programming in ML. In *Proc. ICFP*, pp. 15–27, 2011.

[32] K. Terui. Proof nets and boolean circuits. In *Proc. LICS*, pp. 182–191, 2004.

[33] M. K. Thomsen. A functional language for describing reversible logic. In *Proc. FDL*, pp. 135–142, 2012.

[34] R. Wille, D. Große, *et al.*. RevLib: An online resource for reversible functions and reversible circuits. In *Int'l Symp. on Multi-Valued Logic*, pp. 220–225, 2008.

[35] R. Wille, H. M. Le, G. W. Dueck, and D. Grosse. Quantified synthesis of reversible logic. In *Proc. DATE*, pp. 1015–1020, 2008.

[36] R. Wille, S. Offermann, and R. Drechsler. SyReC: A programming language for synthesis of reversible circuits. In *Forum on Specification Design Languages*, pp. 1–6, 2010.